A Centre-Charge Selection Rule for the Wilson-Line Potential

The fundamental domain of gauge–Higgs unification is representation-dependent

Part III — centre-parity selection, an exact decoupling, and the gate we gave away

Carles Marín*

July 20, 2026

Abstract

Part II gave a gate away. Its first clause — that a representation can display the Standard Model's fractional hypercharges only if its $\mathbb{Z}_4$ centre charge $a+2b+3c$ is odd — was identified as the classical centre-charge congruence and disclaimed: *we claimed none of it.* This paper shows that the disclaimed gate is not a statement about matter at all. It governs the Higgs sector, and for a reason that fits in one line: *advancing the Wilson line by one period is, up to a Weyl reflection, multiplication by the central element $-\mathbf{1} \in Z(SU(4))$ — and the centre charge is by definition how a representation responds to one.* Writing $D(t) = \sum_{\text{weights}} t^{\,n_2-n_4}(-1)^{n_3}$ for the signed generating function of the orbifold twist labels — a projection we *derive* here from the $T^2/\mathbb{Z}_2$ boundary conditions rather than import — the character inherits the centre's response as $D(-t) = (-1)^{a+2b+3c}D(t)$, and hence a dichotomy: a representation pairs its Kaluza–Klein twist multiplicities at *even* m when the centre charge is odd, and at *odd* m when it is even — never both, except on a degenerate class $\mathcal{Z}$ that we classify exactly and certify in Lean. The odd-m case is not ours: it is the hypothesis Akamatsu, Hirose, Maru and Nago verify by inspecting their Table 1 in order to halve the Wilson-line fundamental domain to $\alpha_2 \in [0, \frac{1}{2}]$. Because Part II's gate forces the centre charge *odd* for every representation that can host a quark generation, while the gauge adjoint is unavoidably *even*, the hypothesis inverts precisely on the physical class: we verify directly that the potential of every admissible representation violates period 1 in α_2 by 2–30% of its own scale, so the half-domain reduction does not transfer and the full torus must be scanned. What the admissible class gets instead is an exact decoupling: the notch annihilates the fermion representation's entire contribution to the leading Kaluza–Klein order of the curvature at $\alpha_2 = \frac{1}{2}$, leaving a gauge-only constant identical across all eight admissible representations. The leading order is therefore degenerate *by theorem*, every distinction between candidate representations lives in subleading terms, and a naturalness ranking built on point estimates is fragile for a structural reason rather than a numerical one. We close by replacing the finite-difference Hessian with an integral measure of the basin geometry, validated against it where it is trustworthy and stable where it is not. Finally, the mechanism is neither an accident nor universal: the shift equals a central element only when the Wilson-line-neutral slots carry balanced twist signs, which for a single Wilson-line direction requires N even — so $SU(4)$ is the smallest group that can carry this selection rule at all.

1 The gate we gave away

This paper is about a sentence in its predecessor.

*Independent researcher, karlesmarin@gmail.com. The computations were carried out and cross-checked with Claude (Anthropic) as an AI research assistant against a common machine-verifiable ground truth (exact rational arithmetic on $SU(n)$ weight systems in SageMath, and a `sorry`-free Lean 4 certificate for the central theorem); every number in this paper regenerates from the ancillary scripts.

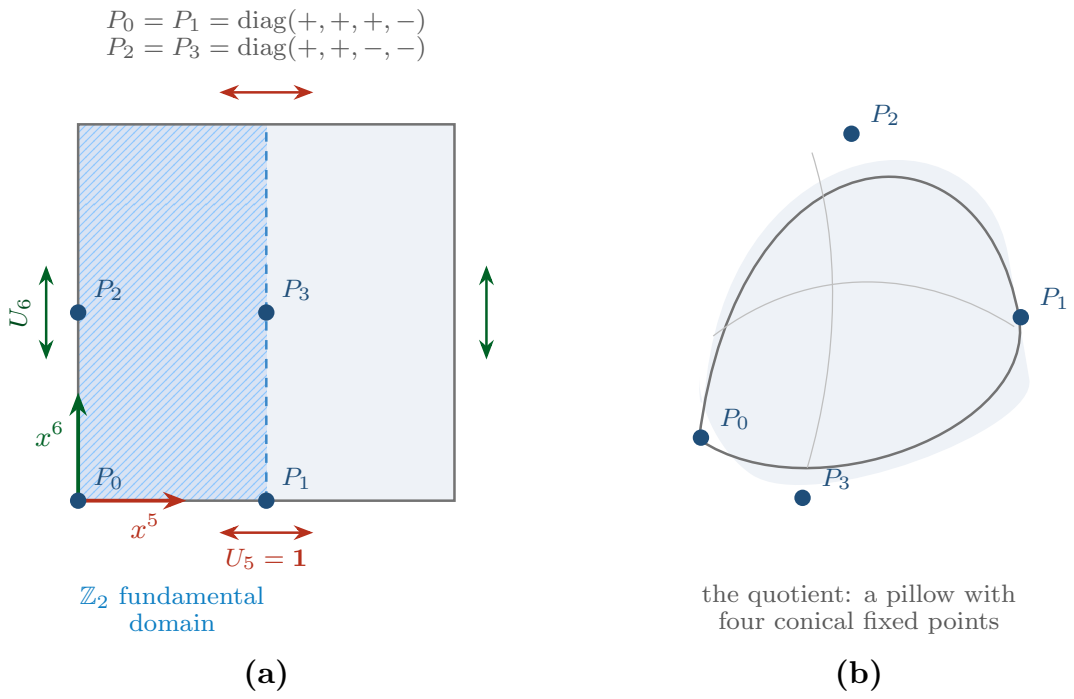


Figure 1: The geometry every later figure lives on. **(a)** The torus T^2 with the $\mathbb{Z}_2$ fundamental domain hatched. Translations around the two cycles are accompanied by U_5 and U_6; inverting AHMN's consistency relations $P_1 = U_5 P_0$, $P_2 = U_6 P_0$ gives $U_5 = \mathbf{1}$ and $U_6 = \mathrm{diag}(1,1,-1,1)$, which is the entire origin of the twist label q (§2). The four $\mathbb{Z}_2$ fixed points carry the parity matrices shown; note $P_0 = P_1$ and $P_2 = P_3$, so only the x^6 direction carries a discrete twist — the asymmetry visible in (2). **(b)** The quotient $T^2/\mathbb{Z}_2$: a pillow with four conical points, one per fixed point. Every heatmap in this paper is the flat picture of (a), drawn over the full period-2 range of the Wilson-line phases.

Part I [1] took a question left open in 2003 — Scrucca, Serone, Silvestrini and Wulzer had observed that on an orbifold the cancellation of anomalies and the vanishing of localized tadpoles are *independent* requirements, and named the missing object an open problem [3] — and met it inside the six-dimensional two-Higgs-doublet $SU(4)$ model of Akamatsu, Hirose, Maru and Nago [4]: a coloured quark block in the dimension-60 representation, a 45-multiplet completion cancelling anomalies and tadpole together, and an Existence theorem showing the tadpole never obstructs. Part II [2] turned the minimality that Part I had merely observed into an exact law: an irreducible (a,b,c) of $SU(4)$ contains a Standard-Model quark cell among its $T^2/\mathbb{Z}_2$ chiral zero modes if and only if

$$(a+2b+3c)\ \text{odd} \quad \wedge \quad b \geq 1 \quad \wedge \quad a+b+c \geq 3. \tag{1}$$

Part II was careful about credit, and nowhere more so than on the first clause. The $\mathbb{Z}_4$ centre charge — Slansky's congruency class $c(R) = a+2b+3c \bmod 4$ [8] — locks fractional hypercharge to a fixed congruence, a fact stated for the Standard Model's own $\mathbb{Z}_6$ by Hucks and by Saller and recast in global-structure language by Tong and by Alonso, Dimakou and West [9, 10, 11, 12]. Part II wrote, in as many words, *"We claim none of it"*, and recorded the clause as *inherited, not invented.*

Giving it away was the right call, and the reason is a long one. The observation that the *centre* of a gauge group carries physics the algebra alone does not goes back to Dirac's quantization condition and was sharpened by 't Hooft's classification of phases by the electric and magnetic charges that survive it [25]: what a theory permits as a line operator, and what fractional charges it can carry, is decided not by $\mathfrak{su}(n)$ but by which quotient of the simply connected group is gauged. Hucks and Saller drew the corresponding conclusion for the Standard Model itself — its true gauge group is $\big(SU(3)\times SU(2)\times U(1)\big)/\mathbb{Z}_6$, and the familiar hypercharges $\frac{1}{6}, \frac{2}{3}, -\frac{1}{3}$ are

precisely the fractions that quotient permits [9, 10]. Tong recast the statement in the language of line operators [11], and Alonso, Dimakou and West asked the experimental question hiding inside it — which fractionally charged states would tell the candidate groups apart [12]. By the time Part II needed the fact, it had been established four separate times in four different languages. One does not claim such a thing.

That was correct, and this paper does not take it back. But it was incomplete in a way we did not suspect, because it left a question unasked. The centre charge entered (1) as an *arithmetic* condition on the matter sector: which weights can carry $\frac{1}{6}, \frac{2}{3}, -\frac{1}{3}$ at all. Nothing in that derivation suggests the same congruence should have anything to say about the *Higgs* sector — about the Wilson-line potential, its symmetries, its vacuum, or the mass it generates. The two live in different parts of the construction and are computed from different data.

So: *does the centre charge know anything about the Higgs potential?* The answer is that it controls it, in a sense we can make exact. We do not claim the congruence. We claim what it does — and what it does is not a restriction on matter but a restriction on the *potential itself.*

Centre-Parity Selection Principle. Let a shift of a Wilson-line modulus be Weyl-equivalent to multiplication by a central element ζ of order d. Then the Fourier support of the one-loop potential splits into d sectors labelled by the centre charge modulo d, and a representation R contributes to a sector only where $\zeta^{\,c(R)} = 1$. The sectors on which $\zeta^{\,c(R)} \neq 1$ are *identically empty* for that representation — not suppressed, absent.

The case proved here is $SU(4)$ on $T^2/\mathbb{Z}_2$, where $\zeta = -\mathbf{1}$ has $d = 2$ and the charge is $c(R) = a + 2b + 3c$: one bit, read in four places — *matter admissibility*, *Kaluza–Klein twist pairing*, the *Fourier support* of the potential, and the *geometry of its vacuum.*

The centre charge has been used for half a century as a condition on which *representations* may carry fractional hypercharge; here it also decides which *harmonics* of the Hosotani potential can exist at all. Figure 2 draws the centre of the group at the centre of the page, with the four sectors it governs around it.

Scope and relation to prior work. As in Part II we separate inherited from new, and again most ingredients are inherited. The centre-charge congruence is classical [9, 10, 11, 12, 8]. The orbifold construction, the twist labels (m, q), the mass equation, the one-loop potential and the fundamental-domain argument are all AHMN's [4]; we quote their Section 4.3 verbatim in §5 because our central physical claim is about the reach of *their* hypothesis, not about an error in it — their argument is correct for their model. The mechanism behind our vanishing theorem is likewise classical: the alphabet $\{1, -1, t, t^{-1}\}$ has determinant -1, so it parametrizes the non-identity component of $O(4)$, where an irreducible representation and its associate $\mu \otimes \det$ have opposite characters and self-associate ones vanish [13]; the branching is Littlewood's restriction rule with the Newell–King modification rules [14, 16, 13], and the determinant technique we use is described by Ayyer and Behrend as routine [22]. What is ours is fourfold: (i) the *dichotomy* (§3) and the exact classification of its degenerate class (§4), machine-certified in Lean 4; (ii) the observation — with direct verification — that AHMN's fundamental-domain reduction *inverts* on the Standard-Model-admissible class (§5); (iii) the *decoupling* corollary and the resulting leading-order degeneracy (§6); and (iv) a derivation of the (m, q) projection from the boundary conditions (§2), which in this construction had been used as a calibrated input. Figure 3 is the route.

Established here: that in this orbifold the unit shift $\alpha_2 \to \alpha_2 + 1$, which is *not* central on its own, is Weyl-equivalent to the central element $-\mathbf{1}$, the U_6 twist supplying the compensating sign — so that the standard centre/n-ality grading of a holonomy potential [41, 42, 43, 44] governs this model, where it is not manifest, and the centre charge becomes a statement about the potential and not only about the matter. *Left open:* all four of its readings, which are the rest of the paper.

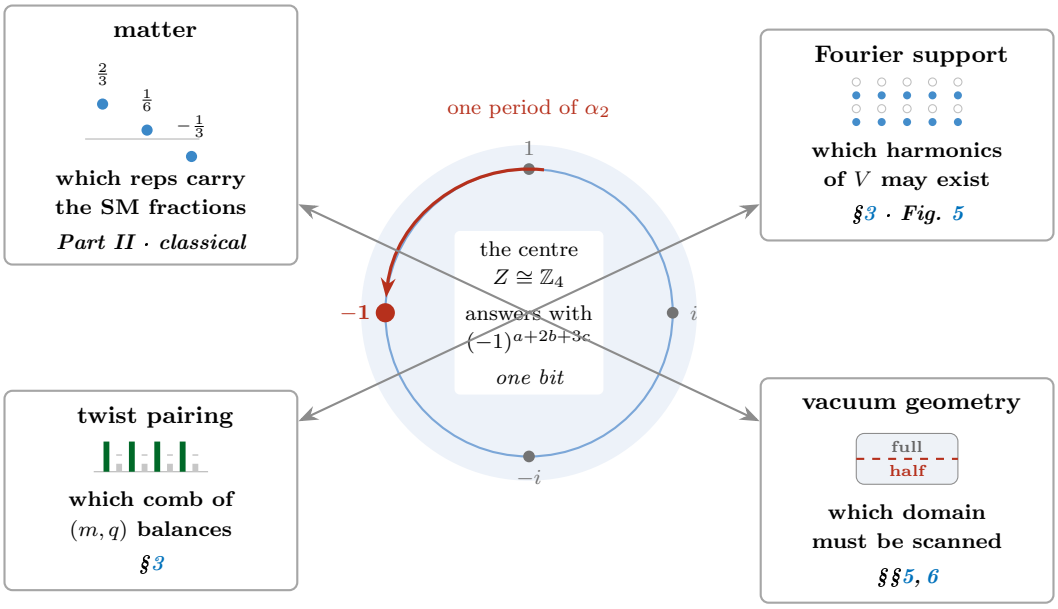


Figure 2: The centre of the group, at the centre of the picture. A period of the Wilson line is a half turn of the $\mathbb{Z}_4$ centre — it lands on $-\mathbf{1}$ — and a representation answers that element with the single bit $(-1)^{a+2b+3c}$. Everything this paper proves is that bit read in four places: which representations may carry the Standard-Model hypercharges (known, and disclaimed in Part II), which comb of twist multiplicities balances, which harmonics of the Wilson-line potential are allowed to be non-zero at all, and how much of the torus must be searched for its vacuum. They look like four separate facts about a model. They are one fact about a group.

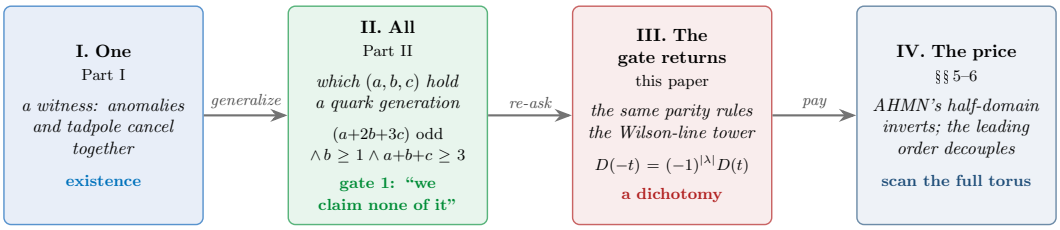


Figure 3: The arc of the series. Part I produced a single consistent model; Part II the exact law of which representations are eligible, explicitly disclaiming its arithmetic gate as classical. This paper asks what that disclaimed gate does elsewhere, and finds it governing the gauge–Higgs potential — with a consequence that must be paid before any vacuum in this class is computed.

2 Where the twist labels come from

Before anything can be proved about the tower, one has to know exactly what labels it.

The labels are old. In 1979 Scherk and Schwarz noticed that a field on a compact space need not come back to itself: it may return acted on by a symmetry, and the resulting twist gives masses that no potential put there [26]. In the same year Manton and Fairlie observed that the extra-dimensional components of a gauge field are scalars in four dimensions, with a quartic coupling fixed by the gauge coupling rather than free [27, 28]; Hosotani then showed that the vacuum of such a scalar is a Wilson line, so that symmetry breaking becomes a statement about holonomy around a cycle [29]. That is the whole idea of gauge–Higgs unification: the Higgs is a twist. Orbifolds entered when Kawamura used a $\mathbb{Z}_2$ projection to split a grand-unified multiplet [30], and the modern constructions — AHMN's included — combine the two, so that a mode carries *both* a Wilson-line phase and a discrete parity. That pair of data is exactly what (m, q) records.

Which is why the pair has to be pinned down before anything else. AHMN attach to each Kaluza–Klein mode a pair of integers (m, q) through twisted boundary conditions [4, Eqs. (3.3)–

(3.4)],

$$f(x^5 + 2\pi R, x^6) = e^{i\pi m\alpha_1} f(x^5, x^6), \qquad f(x^5, x^6 + 2\pi R) = e^{i\pi(m\alpha_2 + q)} f(x^5, x^6), \tag{2}$$

with $\alpha_{1,2} = gRv_{1,2}$ the Wilson-line phases, m "the weight of the generator" and $q \in \{0, 1\}$ periodic or antiperiodic. In terms of the partition coordinates $n = (n_1, n_2, n_3, n_4)$ of a weight this construction is used with $m = n_2 - n_4$ and $q = n_3 \bmod 2$. That assignment is usually justified by checking it against AHMN's Table 1. Checking is not deriving, so we derive it.

The boundary data are AHMN's, and Figure 1 collects them — translations $U_{5,6}$, parities P_i at the four fixed points, and the consistency relations $P_1 = U_5 P_0$, $P_2 = U_6 P_0$ [4, Eq. (2.8)], with the choice [4, Eqs. (2.9)–(2.10)]

$$P_0 = P_1 = \mathrm{diag}(+, +, +, -), \qquad P_2 = P_3 = \mathrm{diag}(+, +, -, -),$$

which breaks $SU(4) \to SU(2)_L \times U(1)_Y \times U(1)_X$. Inverting the relations,

$$U_5 = P_1 P_0^{-1} = \mathbf{1}, \qquad \boxed{U_6 = P_2 P_0^{-1} = \mathrm{diag}(1, 1, -1, 1)} \tag{3}$$

A state of content n in any tensor representation therefore picks up $\prod_i (U_6)_{ii}^{n_i} = (-1)^{n_3}$ under the x^6 translation, which is exactly the periodic/antiperiodic alternative in (2): $q = n_3 \bmod 2$. Two independent checks come for free. First, q is not an extra datum but the *product of the two* $\mathbb{Z}_2$ *parities*: on content n, $P_0 \mapsto (-1)^{n_4}$ and $P_2 \mapsto (-1)^{n_3+n_4}$, whose product is $(-1)^{n_3}$, consistent with $U_6 = P_2 P_0$. Second, and more tellingly, (2) is *asymmetric* — the x^5 condition carries no q and the x^6 condition does. AHMN do not comment on this. It is precisely $U_5 = \mathbf{1}$ against $U_6 \neq \mathbf{1}$: there is no discrete twist along x^5 for a q to describe.

For m: the components of $A_{5,6}$ with four-dimensional massless modes — the $(++)$ entries of [4, Eq. (2.12)] — sit at matrix positions $(1,4), (2,4), (4,1), (4,2)$, which are the two Higgs doublets. The Wilson-line vacuum expectation values lie there, and using $SU(2)_L$ (acting on indices $1, 2$) to align with the second doublet component makes the generator $E_{24} + E_{42}$, conjugate inside the $\{2, 4\}$ block to

$$T = \mathrm{diag}(0, 1, 0, -1), \tag{4}$$

whose weight on content n is $n_2 - n_4$. That a *single* m appears in both halves of (2) is consistent: both cycles carry the same Lie-algebra direction with different magnitudes v_1, v_2. The choice of index pair is $SU(2)_L$ gauge freedom, and it is provably immaterial — the generating function D of §3 is invariant under all 24 assignments $m = n_i - n_j$, $q = n_k \bmod 2$ with i, j, k distinct. The freedom in the derivation and the freedom in the theorem are the same freedom.

Reading AHMN's Table 1 correctly. Applying (T, U_6) to the **35** $= (4, 0, 0)$ gives, for $m \geq 1$, the counts $3, 3, 3, 1, 1, 1, 1$ at $(1,0), (1,1), (2,0), (2,1), (3,0), (3,1), (4,0)$, matching AHMN's Table 1 entry by entry — but as counts of $SU(2)_L$ *multiplets*, not of states. The reason is physical. T does not commute with $SU(2)_L$: moving a box between indices 1 and 2 shifts n_2, hence m, by one. A multiplet of size s therefore does not carry a single m; it spreads over s *consecutive* values, contributing exactly one state to each. We verify this on the **35** with no exceptions, and the identity "#weights at (m, q) = #multiplets spanning (m, q)" holds exactly. That the Wilson line splits $SU(2)_L$ multiplets is not bookkeeping — it is electroweak symmetry breaking, which is what the Wilson line is for.

Established here: q is the $U_6 = P_2 P_0^{-1}$ twist, equivalently the product of the two orbifold parities, and m is the weight under the Wilson-line generator; AHMN's Table 1 is reproduced *and* explained. *Left open:* nothing at this level — what remains an input is AHMN's parity choice itself, which is their model definition, as it should be.

3 Centre-parity selection, and the dichotomy it forces

With the labels fixed we can ask a sharp question about the tower — and the way to ask it is itself borrowed.

There is a long habit in representation theory of learning about a character by evaluating it at a *distinguished* group element rather than studying it in general. Weyl's *Classical Groups* is built on the technique [32], and Littlewood turned it into a machine: feed a Schur function an alphabet with structure — variables in ± pairs, say — and the answer is not merely simpler but combinatorial, vanishing outright unless the partition passes a divisibility test [33]. The habit is very much alive: Ciucu and Krattenthaler found that characters at reciprocal pairs x, x^{-1} factorize [36], Ayyer and Behrend widened the class [22], and Ayyer and Kumari classified exactly which partitions survive a root-of-unity twist [20, 21]. In every case the special element converts a question about representations into arithmetic on a Young diagram.

Our tower hands us such an element for free. The two orbifold parities and the Wilson line together evaluate the $SU(4)$ character at an alphabet that is half ± 1 and half reciprocal — and, as we will see, that alphabet is not a convenient choice but the one the physics dictates. So we follow the habit. Define the signed generating function of the twist labels,

$$D(t) \;=\; \sum_{\text{weights}} t^{\,n_2-n_4}\,(-1)^{n_3} \;=\; s_\lambda\big(1,-1,t,t^{-1}\big)\,, \qquad \lambda=(a{+}b{+}c,\ b{+}c,\ c,\ 0), \tag{5}$$

the second equality because $s_\lambda(x)=\sum_{\text{weights}}\prod_i x_i^{n_i}$ evaluated at $x=(1,t,-1,t^{-1})$, reordered by symmetry. Since the (m,q) classes partition the weights and $(-1)^{n_3}$ is constant on each,

$$[t^m]\,D(t) \;=\; \mathrm{mult}(m,0)-\mathrm{mult}(m,1), \tag{6}$$

so D records exactly the failure of the two twist sectors to balance, level by level. Call (a,b,c) *even-notched* if $\mathrm{mult}(m,0)=\mathrm{mult}(m,1)$ for every even m, and *odd-notched* if the same holds for every odd m. The first happens iff D is an odd Laurent polynomial, the second iff D is even.

How this was found, since it matters for how much to trust it. The lemma below is two lines, which makes it look like the sort of thing one writes down first and then goes looking for an application. That is not what happened, and the actual order is worth recording because it is the same order Part I recommended: notice a regularity, then ask whether it is a law.

The even-m notch appeared first as an *empirical* oddity. Tabulating (m,q) multiplicities across the admissible catalogue while working on something else, we noticed that every admissible representation balanced its even-m levels exactly, and that AHMN's **35** — the one representation in the table that cannot host a quark generation — did not, missing by 3 against 1 at $m=2$. Eight cases and one exception is a pattern, not a theorem, and our first explanation for it was wrong: we conjectured the responsible clause was Part II's second gate, $b\geq 1$, and falsified it within the hour on the explicit counterexamples $(2,1,0)$ and $(0,2,0)$, which have $b\geq 1$ and no notch. Only after that failure did we look for the invariant that the parity of the centre charge supplies, and only then did the two-line proof exist to be written. The classification of the degenerate class in §4 came later still, forced by the cases the lemma leaves open. We record the sequence because a reader is entitled to know whether a clean statement was reverse-engineered from a clean proof or extracted from messy data; this one came from the data, and the counterexamples that killed the first attempt are in the ancillary scripts alongside the ones that confirm the second.

Lemma 1 (Parity)**.** *$D(-t)=(-1)^{|\lambda|}D(t)$, with $|\lambda|=a+2b+3c$ the number of boxes of λ.*

Proof. As multisets $\{1,-t,-1,-t^{-1}\}=-\{1,t,-1,t^{-1}\}$; s_λ is symmetric and homogeneous of degree $|\lambda|$. □

The principle: a unit shift of the Wilson line is a central gauge transformation. Written that way the lemma looks like an accident of Schur functions, and it is worth saying plainly why it is not — because the reason is the whole content of this paper.

The generating variable tracks the twist level, $t^m \leftrightarrow e^{i\pi m\alpha_2}$, so $t \to -t$ is nothing other than $\alpha_2 \to \alpha_2 + 1$: one unit of Wilson line. Write $g(\alpha_2)$ for the group element the orbifold hands us — the U_6 twist times the Wilson line, with eigenvalues $(1,\ t,\ -1,\ t^{-1})$ in the slots of §2. Then

$$\begin{aligned} g(\alpha_2{+}1) &= \mathrm{diag}\big(1,\ -t,\ -1,\ -t^{-1}\big), \\ (-\mathbf{1})\cdot g(\alpha_2) &= \mathrm{diag}\big(-1,\ -t,\ \ 1,\ -t^{-1}\big). \end{aligned}$$

These are *not* the same matrix, and we are careful not to say they are: the shift on its own, $\mathrm{diag}(1,-1,1,-1)$, is not central. They agree *up to a Weyl element*, the unique permutation relating them being the transposition of slots 1 and 3. That is enough, because a character is a class function:

$$\boxed{g(\alpha_2{+}1) \ \sim_W\ (-\mathbf{1})\cdot g(\alpha_2), \qquad -\mathbf{1} \in Z\big(SU(4)\big) \cong \mathbb{Z}_4.} \tag{7}$$

The centre of $SU(4)$ is $\{\zeta\mathbf{1} : \zeta^4 = 1\}$, and $-\mathbf{1}$ is its unique element of order two ($\det(-\mathbf{1}) = (-1)^4 = 1$, so it does lie in $SU(4)$). On an irreducible representation $\zeta\mathbf{1}$ acts by the scalar $\zeta^{|\lambda|}$ — that is what the n-ality *is* — so $-\mathbf{1}$ acts by $(-1)^{|\lambda|}$, and $\chi_\lambda(g(\alpha_2{+}1)) = (-1)^{|\lambda|}\chi_\lambda(g(\alpha_2))$. That is the parity lemma, and (7) is why it had to hold:

> *Advancing the Wilson line by one period is, up to a Weyl reflection, a central gauge transformation — and the centre charge is by definition how a representation responds to one.*

The Weyl element is not a technicality to be apologised for; it is where the orbifold enters. The permutation that does the job exchanges the *free* slot, where the alphabet is $+1$, with the slot carrying the U_6 twist, where it is -1. Without that twist there would be no -1 in the alphabet to pair with the $+1$, the negation could not be undone by any permutation, and no selection rule would exist. The same U_6 that produces the label q in §2 is what makes the centre charge visible to the potential at all.

Two consequences are worth stating before the details. First, nothing here is special to Schur functions: in any construction where a shift of a modulus realizes a central element up to Weyl, a centre-charge selection rule follows, whether or not anyone has written it down. Second, the *mechanism* is not special to $SU(4)$ either — every $SU(N)$ has a centre $\mathbb{Z}_N$ — but the *congruence* that appears is: it is fixed by the order of the central element the shift realizes, which is two here and need not be elsewhere. §10 takes that up, since it is the first question a reader should ask.

The lemma itself is elementary, and we present it as bookkeeping rather than as a result. Its consequence is not.

Corollary 1 (Dichotomy)**.** *Every irreducible representation of $SU(4)$ is even-notched or odd-notched, and not both, except on the degenerate class $D \equiv 0$ where it is both:*

$$|\lambda| \text{ odd} \implies \text{pairing at } \textbf{even}\ m, \qquad |\lambda| \text{ even} \implies \text{pairing at } \textbf{odd}\ m.$$

Figure 4 is the corollary with the algebra taken out. Two remarks fix the meaning. First, $|\lambda| = a + 2b + 3c \equiv a + c \pmod 2$, so the parity in question is Part II's gate 1 exactly — and it is the $\mathbb{Z}_2$ *quotient* of the $\mathbb{Z}_4$ centre charge, not the centre charge itself. Second, the same parity was already doing silent work in Part II: the chiral zero-mode count $N = (b+1)(a+c+1)/2$ is an integer only because $a + c$ is odd whenever the centre charge is odd, a parenthetical remark there. Three appearances, in three sectors that share no coordinate: the admissibility gate, the integrality of the zero-mode count, and now the tower's twist balance. It is the invariant doing the work, not the frame.

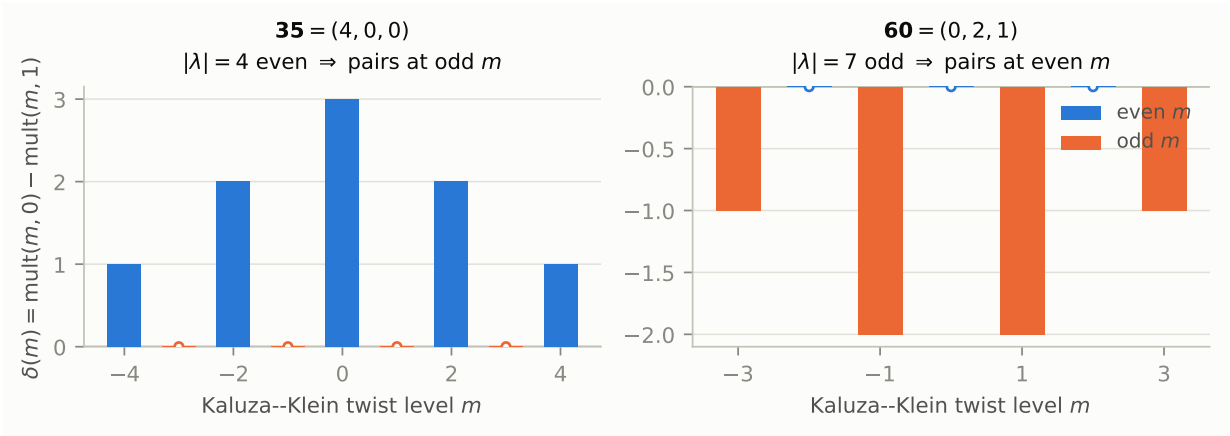


Figure 4: The dichotomy, seen directly. The signed imbalance $\delta(m) = \mathrm{mult}(m,0) - \mathrm{mult}(m,1)$ for AHMN's fermion (left) and for the smallest representation able to host a quark generation (right). Each is a *comb*: one vanishes on every even tooth, the other on every odd one, and which comb a representation gets is decided by a single bit, the parity of its centre charge. Hollow markers sit on the teeth that vanish identically. No representation has both combs full, and outside the class $\mathcal{Z}$ of §4 none has both empty.

The notch is a hole, not a cancellation. It is worth seeing what Corollary 1 does to the potential itself, because the statement is more concrete than "multiplicities balance". Expanding V in Fourier modes on the torus, the amplitude at the frequency vector $m\,(k_1, k_2)$ — the m-th harmonic along the direction (k_1, k_2) — carries the factor $\sum_q \mathrm{mult}(m,q)\,(-1)^{k_2 q}$, which is the *total* multiplicity when k_2 is even and exactly $\delta(m)$ when k_2 is odd. The notch therefore does not arrange a delicate cancellation among contributions: it empties an entire sublattice of Fourier space. For a representation of odd centre charge, *every* mode with k_2 odd and m even is identically absent from the potential — not small, absent. Figure 5 shows the eight admissible representations leaving precisely the same seats vacant, and AHMN's **35** filling all 95 of them.

Established here: exactly one of the two pairings holds, which one is decided by the centre charge's parity, and the consequence is an identically empty sublattice of the potential's Fourier support. *Left open:* the degenerate class where both hold, and — the point of the paper — what the odd-m case is doing in the literature already.

4 The degenerate class

Corollary 1 is silent when D vanishes identically. That class is small, exactly computable — and its existence is the fingerprint of a group that is not connected.

$O(n)$ has two components, and everything awkward about its representation theory descends from that. Its irreducible representations come in *associate* pairs μ and $\mu \otimes \det$, indistinguishable on the rotations and opposite on the reflections; some are their own partner. Because Young diagrams were built for GL and S_n, labelling $O(n)$ representations by them produces "non-standard" labels that must be repaired, and Newell wrote the repair rules in 1951 [16], with King systematizing them for all three classical families two decades later [14]. The awkwardness is not bookkeeping fussiness; it is the disconnectedness showing through the notation. Koike and Terada's influential reformulation resolves it in the cleanest possible way — by working in $R(SO(n))$ and treating the $O(n)$ labels as a basis [18] — which is elegant, universally used, and, for our purposes, precisely the wrong move: identifying μ with $\mu \otimes \det$ is exactly the information we need.

That is the sense in which the class below exists. It consists of the representations whose branching is *blind* to the reflection — self-associate through and through, so that every contribution cancels against a partner. The physics saw it first, as a tower with both combs empty.

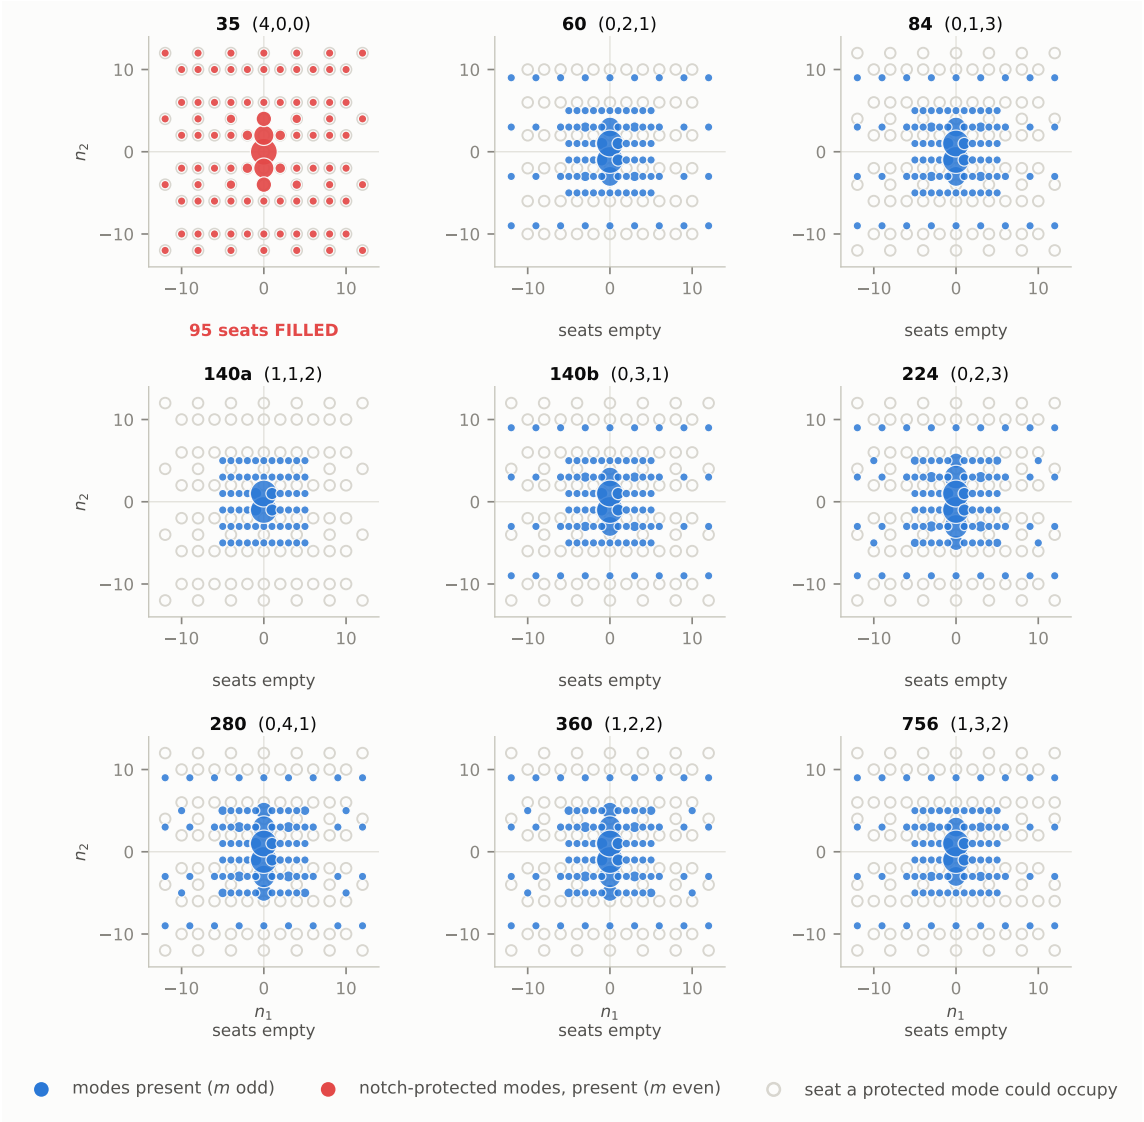


Figure 5: The notch, seen as absence. Each panel is the Fourier plane of the one-loop potential for one representation, restricted to the odd-k_2 shells — the sector whose amplitude is $\delta(m)$. Filled discs are modes the potential actually carries, sized by amplitude; hollow rings mark every seat a *notch-protected* mode (m even) could occupy. Eight representations leave all of those seats empty, and they are exactly the eight that can host a Standard-Model quark generation. The ninth, AHMN's **35**, fills them. Nothing was fitted or thresholded: the empty seats are exact zeros, and the pattern is the single bit $a+2b+3c$ mod 2 made visible.

Theorem 1 (Degenerate class)**.** $D \equiv 0$ *if and only if*

$$\boxed{b \text{ odd } \wedge \ \big[(a, c \text{ both odd}) \ \vee \ (a = c \text{ even})\big]}$$

Consequently the even-m notch holds iff $a+2b+3c$ is odd or $(a, b, c) \in \mathcal{Z}$, and since every member of $\mathcal{Z}$ has $|\lambda|$ even, $\mathcal{Z}$ is disjoint from the class satisfying Part II's gate.

The two branches are not on the same footing and we separate them. The first, a, b, c all odd, says every μ_j has the same parity, so the two constant rows of the bialternant become proportional; it is the classical β-set parity obstruction and it is not special to our alphabet at all — it is exactly the vanishing locus of $s_\lambda(1, -1, y_1, y_2)$ with y_1, y_2 *free*. It should be read as standard. The second, $a = c$ with b odd, is not a symmetric-function identity in free variables: released from $y_2 = y_1^{-1}$ it stops vanishing, and over the labels up to 5 the two loci differ in exactly the six shapes $(0, 1, 0)$, $(0, 3, 0)$, $(2, 1, 2)$, $(2, 3, 2)$, $(4, 1, 4)$, $(4, 3, 4)$, all of this branch. It lives on the orthogonal coset and nowhere else, which is why its natural statement is the one of §10: the $GL(4)$ character vanishes identically on $O(4)^-$, equivalently $\operatorname{Res} V_\lambda \cong \operatorname{Res} V_\lambda \otimes \det$.

One ancestor deserves naming for this branch specifically. Ayyer and Behrend's factorization theorem [22, Thm. 1] is stated for partitions with $\lambda_1 + \lambda_4 = \lambda_2 + \lambda_3 = 2k$, which for four rows is exactly $a = c$, and its half-partition case is exactly b odd — their shape hypothesis is our branch, in its more delicate sub-case. Their alphabet is $(x_1, x_2, \bar{x}_1, \bar{x}_2)$, the identity component, so the theorem does not reach ours, and their prefactor $\prod(x_i^{1/2} + \bar{x}_i^{1/2})$ degenerates precisely at $x_i = -1$. We claim no more than that the family is theirs and the coset is ours.

Proof sketch. Write $D = N/V$ with $N = \det(y_i^{\mu_j})$ the bialternant on the ordered alphabet $y = (1, -1, t, t^{-1})$ and $\mu = (a{+}b{+}c{+}3,\ b{+}c{+}2,\ c{+}1,\ 0)$. Laplace-expand N along the two *constant* rows: the 2×2 minor on columns $j < k$ is $(-1)^{\mu_k} - (-1)^{\mu_j}$, so only columns of opposite μ-parity survive, while every complementary minor on the rows (t, t^{-1}) equals $[\mu_l - \mu_m]$ with $[x] := t^x - t^{-x}$. The eight parity classes of (a, b, c) fall into four types. If all μ_j are even — that is, a, b, c all odd — every surviving minor vanishes and $N \equiv 0$. If exactly one μ_j has the minority parity, N is a combination of three brackets whose largest argument is the sum of the other two, hence cannot cancel. Of the three classes with a two-two parity split, two contain the bracket $[P]$ with $P = a{+}b{+}c{+}3$ strictly larger than every other argument, so again no cancellation; the third, a, c even and b odd, gives $N/2 = [Q] - [R] + [P{-}Q] - [P{-}R]$ with all four arguments positive, which vanishes iff $\{Q, P{-}Q\} = \{R, P{-}R\}$ as multisets, and since $Q = R$ would force $b = -1$ this happens iff $a = c$. The only analytic input is that $\{[x]\}_{x>0}$ is linearly independent, which is the statement that distinct monomials are independent. □

The theorem is verified four independent ways — Weyl weight multiplicities, an exact bialternant sweep over 10 648 representations to $(a, b, c) \le 21$ with no mismatch, symbolic checks of each case formula, and an adversarial audit — and certified in Lean 4 (Table 1). Figure 6 puts the theorem and Part II's criterion on the same plate: $\mathcal{Z}$ occupies its own cells and never meets a representation able to host a quark generation, which is Corollary 1's disjointness seen rather than deduced.

What is new here, and what is not. The mechanism is classical and we state it plainly. The alphabet $\{1, -1, t, t^{-1}\}$ has determinant -1, so it is the generic semisimple element of the non-identity (*improper*) component $O(4)^-$; an $O(4)$ irreducible representation and its associate $\mu \otimes \det$ have opposite characters there, and every $\mu = (\mu_1, \mu_2)$ with $\mu_2 \ge 1$ is self-associate and contributes nothing [13, Thm. 5]. Hence $D \equiv 0$ iff the Littlewood branching multiset of λ is stable under the associate involution, i.e. iff $c^\lambda_{(m)} = c^\lambda_{(m,1,1)}$ for all $m \ge 1$ and $c^\lambda_\varnothing = c^\lambda_{(1^4)}$. Our proof *is* that computation in another basis: one checks $V = -2[1]^3$, so the brackets $[x]$ are the non-identity-component characters $\chi^-_{(x-1)}$ up to normalisation. We give the bialternant route because it avoids the modification bookkeeping — Littlewood's rule is unmodified only

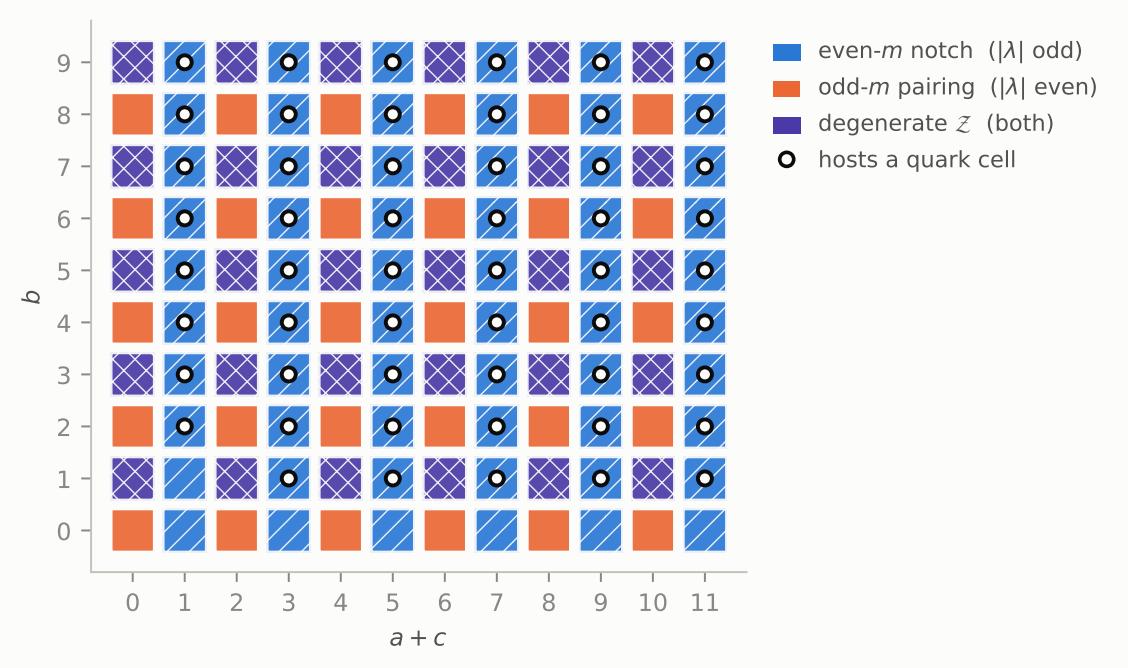


Figure 6: The whole landscape on one plate. Each cell is a label pair $(a+c,\, b)$ — the two coordinates the classification actually depends on — coloured by which pairing the tower satisfies. The blue column-parity is Part II's gate 1 in disguise: those are precisely the representations that can host a quark generation, ringed where the remaining two clauses ($b \geq 1$, $a+b+c \geq 3$) also hold. Violet cells are the degenerate class $\mathcal{Z}$ of Theorem 1, which satisfies both pairings and, as the figure makes visible, never touches a ring. Everything in this paper is a statement about the boundary between the blue and the orange.

for $\ell(\lambda) \leq n/2 = 2$, i.e. only for $c = 0$, and outside that range King's modification rules are required [14, 16, 13], with illegal $O(4)$ labels appearing with multiplicity already for labels as small as $(2, 1, 2)$. That is a practical advantage, not a conceptual one. What we claim is the explicit arithmetic classification in Dynkin labels and its Lean certificate, not the phenomenon and not the technique. The character value at an involution $(1^p, (-1)^q)$, which our bialternant route reproduces along the way, is likewise not ours: it is Karmakar [23, Thm. 4.1], with the same parity-split-of-$\lambda + \rho$ device and the same column-subtraction-then-Laplace proof.

Established here: $\mathcal{Z}$ exactly, disjoint from Part II's admissible class.

Four questions this paper cannot answer. Theorem 1 tells us exactly *when* the shadow character vanishes, and the rest of this paper spends that fact. It leaves four things standing, and we state them as questions rather than as regrets because they are answerable and because the answers turned out to be one object:

Q1 Is $\mathcal{Z}$ a corollary of some general criterion for a $GL(n)$ character to vanish on the non-identity component, or is it particular to $SU(4)$?

Q2 The theorem says where D is zero. *What is D when it is not?* We use its coefficients throughout and never learn what function they belong to.

Q3 Why exactly two branches? Figure 6 shows no third cause of vanishing, and we have no reason why there could not be one.

Q4 The moments $\sum_m m^{2r}\delta(m)$ decide everything that survives §6, and we obtain them by enumerating the weight system of each irreducible. Is there structure there, or is enumeration the only route?

Answered elsewhere (v4). All four are settled in Part IV, and by a single statement: a closed form for $s_\lambda(1, -1, t, t^{-1})$ as $\pm$ a product of exactly three $SU(2)$ characters read off the 2-quotient

Lean theorem (`NotchCentreCharge.lean`)	content
`notch_degenerate_iff`	$N\,a\,b\,c = 0 \leftrightarrow (\mathrm{Odd}\,b \wedge ((\mathrm{Odd}\,a \wedge \mathrm{Odd}\,c) \vee a = c))$ — Theorem 1, over $\mathbb{Z}[T;T^{-1}]$, no `decide` on the determinant
`notch_parity`	Lemma 1 as a division-free identity on the bialternant
`degenerate_centre_charge_even`	every member of $\mathcal{Z}$ has $a+2b+3c$ even
`degenerate_disjoint_admissible`	$\lvert\lambda\rvert$ odd $\Rightarrow N \neq 0$; the two classes never meet

Table 1: The machine-checked core. All statements are `sorry`-free and depend only on `propext`, `Classical.choice`, `Quot.sound`. The linear-independence step is replaced by extraction of a single Laurent coefficient at the strictly largest bracket argument, which is finite and effective.

of λ, or zero. **Q1** yes, and by something stronger than a criterion; **Q2** that product; **Q3** because a product of three characters has exactly two ways to be zero; **Q4** there is structure — δ is a threefold convolution, so every moment is a polynomial in the three indices, and the enumeration is unnecessary. That is a statement about symmetric functions, and it sits beside the factorisation programme of Ciucu–Krattenthaler and Ayyer–Behrend [36, 22] without belonging to it or to the root-of-unity line. Part IV develops it; what it hands back to the physics is the subject of the next paragraph.

It is worth saying what comes back the other way, because the exchange is not one-sided. The closed form makes D a product of exactly three $SU(2)$ characters, so $\delta(m)$ is a threefold convolution and its moments $\sum_m m^{2r}\delta(m)$ are *polynomial* in the three indices. Those moments are exactly the quantities this paper had to extract by enumerating the weight system of each irreducible, and they are what decides the comparison between admissible representations once §6 has removed the leading term. The question we exported to Part IV as a curiosity about symmetric functions returns as the analytic handle on the subdominant structure we could describe here and not compute.

5 The price: a fundamental domain that does not transfer

Here is the turn. The odd-m half of Corollary 1 is not a curiosity of ours. It is a hypothesis already in use — and it is load-bearing.

In computing the one-loop Wilson-line potential, AHMN reduce the search region. They observe [4, §4.3]:

> *"For even m, the combination $(m\alpha_2+q)/2$ is invariant under $\alpha_2 \to \alpha_2+1$ up to an integer shift, and hence each such tower is periodic with period 1 in the α_2 direction. For odd m, the spectrum with $(m,q)=(m,1)$ at a given α_2 is equivalent to that with $(m,q)=(m,0)$ at α_2+1. Hence, if the effective potential contains an equal number of KK mode with $(m,0)$ and $(m,1)$ for odd m, their combined contribution is also periodic under $\alpha_2 \to \alpha_2+1$."*

We quote the even-m sentence as well, because it works for the argument rather than against it: the even-m towers are individually period-1 whatever the representation, so the odd-m sector is the *only* place where the periodicity can be obstructed — which is why a representation that fails there loses the reduction outright. They then discharge the hypothesis by inspection — *"From Table 1, we see that..."* — for their own field content, and obtain the fundamental region [4, Eq. (4.15)]

$$0 \le \alpha_1 < 1, \qquad 0 \le \alpha_2 \le \tfrac{1}{2}. \tag{8}$$

sector	(a,b,c)	$\lvert\lambda\rvert$	odd-m pairing (AHMN's hypothesis)	even-m notch
gauge adjoint **15**	$(1,0,1)$	4 even	yes	no
AHMN fermion **35**	$(4,0,0)$	4 even	yes	no
every admissible rep	gate 1	**odd**	**no**	**yes**

Table 2: The inversion. Part II's gate 1 forces $|\lambda|$ odd for any representation able to host a quark generation, while the gauge adjoint is unavoidably even — so the two sectors land on opposite parities and no admissible completion can restore the odd-m pairing that (8) needs.

	hypothesis	status of that hypothesis	conclusion drawn
AHMN [4] §4.3	equal multiplicities of $(m,0)$ and $(m,1)$ at **odd** m	checked by inspecting their Table 1, for their own field content	$\alpha_2 \in [0,\frac{1}{2}]$ (8)
this paper, §3	none — the pairing is *derived*	a theorem: which comb a representation gets is $(a{+}2b{+}3c)$ mod 2	the dichotomy, Cor. 1
consequence	Part II forces $\lvert\lambda\rvert$ odd on every admissible rep; the gauge adjoint is even	so the odd-m pairing fails on the whole physical class, verified on V itself	$\alpha_2 \in [0,1]$: the full torus, Cor. 2

Table 3: The relation to the result this paper builds on. AHMN's argument is correct and their check is valid for their model; what was not available at the time is that their hypothesis is decided by the centre charge, and therefore inverts exactly on the representations able to host the Standard Model.

Their check is correct. But the hypothesis is not a property of their model; by Corollary 1 it is a property of the *representation*, and it flips with the centre charge. Table 2 is the whole argument.[1]

For the total vacuum-expectation-value-dependent spectrum, gauge plus fermion, the precondition of (8) holds for the **35** and fails for every one of the eight admissible representations of dimension ≤ 756.

Table 3 sets the two arguments side by side, because the relation between them is easy to misread. We do not contradict a result; we identify the hypothesis under it, show that the hypothesis is a property of the representation rather than of the model, and observe that it inverts on the class one would want to use next.

The conclusion, tested rather than inferred. AHMN prove *pairing* ⇒ *period* 1. Showing that the pairing fails does not by itself show that the periodicity fails — that would be denying the antecedent, and some other route to the same periodicity could exist. So we test the conclusion directly on the potential. Evaluating V over the period-2 torus by fast Fourier transform (it is exactly a two-dimensional cosine series, since $k_2 q$ is an integer and the sine term dies) and comparing $V(\alpha_1,\alpha_2{+}1)$ with $V(\alpha_1,\alpha_2)$:

rep	**35**	**60**	**84**	**140a**	**140b**	**224**	**280**	**360**	**756**
$\max\lvert V(\alpha_2{+}1)-V\rvert$/scale	**0**	0.30	0.21	0.045	0.12	0.19	0.12	0.095	0.023

The **35** is period-1 *exactly*; every admissible representation violates it by 2–30% of the potential's own scale — a physical amount, not numerical noise (Figure 7). The reflection $\alpha_2 \to -\alpha_2$

[1] That the period depends on the representation is not invisible in the earlier literature so much as unremarked in it: the general five-dimensional formula of Haba, Takenaga and Yamashita [38] already carries adjoint-type contributions as $\cos(2\pi n a)$ against fundamental-type ones as $\cos(\pi n a)$, periods differing by exactly the factor of two at issue. The distinction is present in their equations and is drawn as a consequence of the mode sums; we have found it stated nowhere as a rule about which harmonics may exist, and that is the gap the centre charge fills.

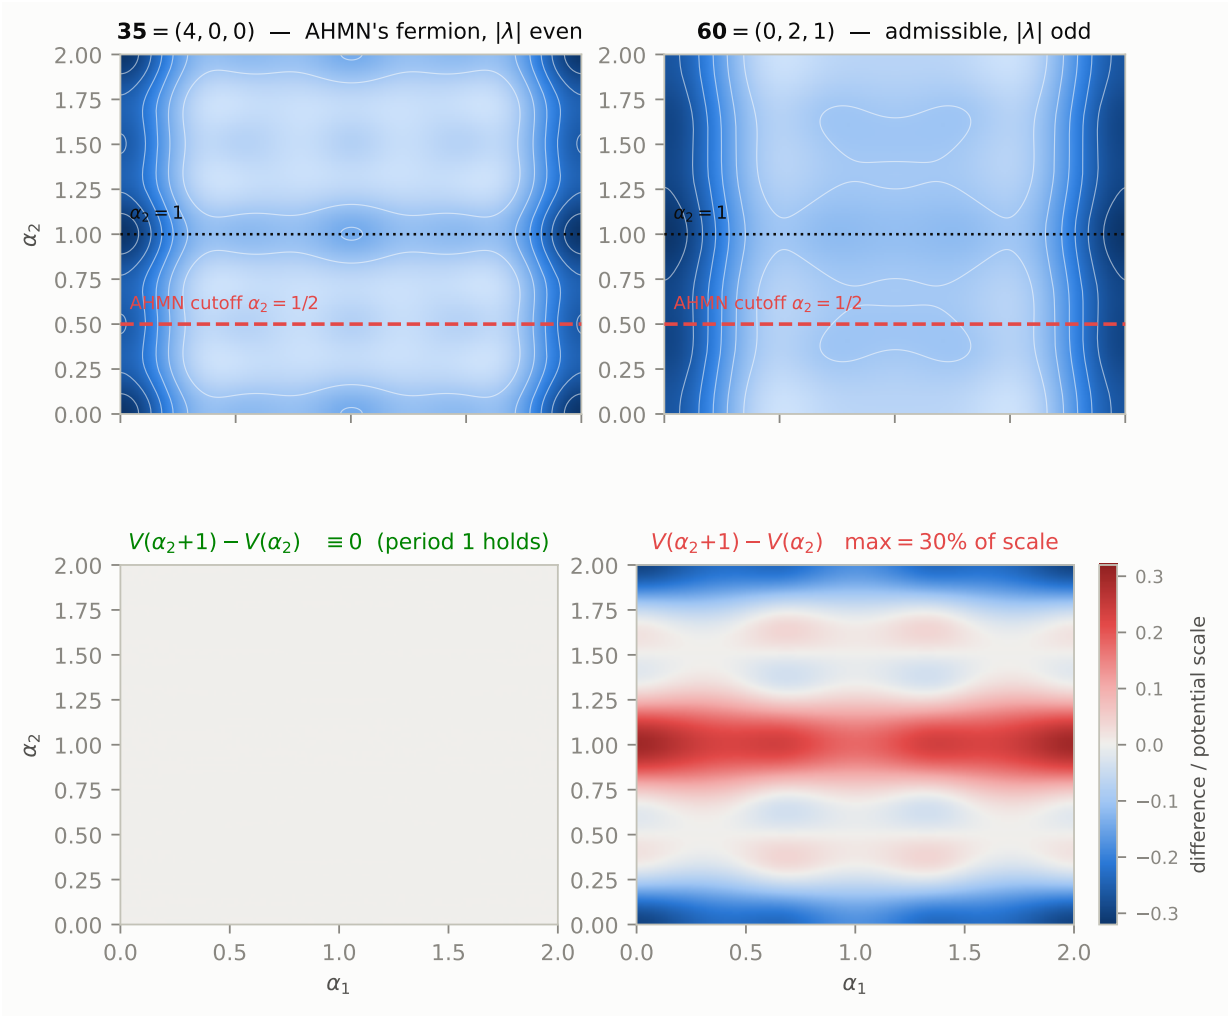


Figure 7: The central claim, without algebra. **Top:** the one-loop Wilson-line potential over the full period-2 torus, for AHMN's fermion (left) and for the smallest admissible representation (right); the dashed line is the edge of the fundamental region (8) that the literature uses. **Bottom:** the direct test, $V(\alpha_1, \alpha_2{+}1) - V(\alpha_1, \alpha_2)$ on the same scale. On the left it is *identically zero* — flat grey, the periodicity holds and the half-domain is legal. On the right it is a structured 30% of the potential's own amplitude. Nothing about the right-hand panel is smaller, subtler, or more numerical than the left; the symmetry is simply absent, and every candidate representation that can host the Standard Model is on that side.

does survive in all nine. Hence for admissible representations the fundamental region is $\alpha_2 \in [0, 1]$: exactly twice AHMN's.

Corollary 2. *The reduction* (8) *does not transfer from AHMN's model to any Standard-Model-admissible representation. A minimization for such a representation must scan the full torus; a search restricted to the half-domain can return a boundary point that is not a vacuum.*

Relation to the boundary-condition equivalence-class programme. There is a developed literature on gauge redundancy in exactly this setting, and it is worth saying precisely why it does not contain the statement above. Kawamura and collaborators, and more recently Takeuchi and Inagaki, classify *equivalence classes of boundary conditions* on $S^1/\mathbb{Z}_2$ and $T^2/\mathbb{Z}_m$ under the gauge transformations that connect them [37, 45, 46, 47]. Their invariant is a trace conservation law: $\mathrm{tr}\, P_i$ is preserved at each fixed point, so the eigenvalue degeneracies of the boundary-condition matrices cannot change within a class.

That is a different invariant on a different object. Theirs is evaluated on the $N \times N$ *boundary-condition matrices* and constrains which discrete boundary conditions may be identified; ours is

evaluated on an arbitrary *irreducible representation* at fixed boundary conditions, and constrains which harmonics of the potential may be non-zero. The two do not compute one another.

One equation of that programme does deserve to be named as an antecedent, and we name it. Haba, Hosotani and Kawamura exhibit, for $SU(2)$ and as an illustration that boundary conditions move under gauge transformations, the equivalence $(\tau_3, \tau_3, \mathbf{1}) \sim (\tau_3, -\tau_3, -\mathbf{1})$ [37, Eq. (2.20)], obtained from $\Omega = \exp\{i(\alpha y/2\pi R)\tau_2\}$ at $\alpha = \pi$ — which in the normalization used here is exactly one period of Wilson line. That is the $N = 2$ shadow of (7), and a reader who knows it should be told why it does not already contain the statement. In $SU(2)$ the block *is* the group, so the holonomy $-\mathbf{1}$ is central for free and no compensation is needed. Embedded in $SU(N)$ along their own relation [37, Eq. (2.21)] the same matrix becomes $\mathrm{diag}(\dots, -1, -1, \dots)$, which is not central, and no centre statement follows. What makes (7) work at $N = 4$ is the twist, not the shift; that is the step with no counterpart there. Consistently, the words *centre*, *Weyl*, *period* and *n-ality* do not occur in [37], and its effective potential is computed only at vanishing Wilson-line phases. Two features of that programme make the separation sharp rather than a matter of emphasis: it declines to compute an effective potential at all ("we do not carry them out since our purpose is not to study the dynamics of models but to classify the BCs" [45, p. 7]), and it notes that the trace of the *translation* matrix is not in general conserved [46, fn. 3] — the translation sector being precisely where our shift $\alpha_2 \to \alpha_2 + 1$ lives.

And no other identification rescues it. Enlarging a fundamental domain invites the obvious objection: perhaps some *other* identification, not considered here, shrinks it again. On a torus the natural candidates are the modular transformations, which act on the Wilson-line phases by mixing the two cycles — the same group that organizes the modular flavour-symmetry programme [48, 49]. We therefore measured, rather than assumed, which lattice transformations V actually respects. The surviving group is exactly $\{\alpha_1 \to \pm\alpha_1\} \times \{\alpha_2 \to \pm\alpha_2\}$ together with the translations: the reflections and $S^2 = -\mathbf{1}$ hold to machine zero for every representation tested, while S itself, the exchange $\alpha_1 \leftrightarrow \alpha_2$, and the shear $\alpha_1 \to \alpha_1 + \alpha_2$ are all broken at order unity. This is what one should expect from (3): $U_5 = \mathbf{1}$ while $U_6 \neq \mathbf{1}$, so the two cycles are inequivalent and nothing that mixes them can survive the orbifold projection. The consequence for us is that $\alpha_2 \in [0, 1]$ is not merely larger than (8) — it is the correct and *complete* fundamental region, with no further identification available to reduce it.

One residual is worth flagging rather than interpreting. The shear $\alpha_2 \to \alpha_1 + \alpha_2$ is broken, but by a margin that falls monotonically with the dimension of the representation: 0.24, 0.18, 0.13, 0.075 and 0.027 for the **35**, **60**, **84**, **280** and **756** respectively. Whether that is an asymptotic recovery of a modular symmetry in the large-representation limit, or merely the smoothing of the potential as more modes contribute, we do not know; a single monotone sequence is not evidence, and we state it as an observation only.

This is not an erratum for [4], whose field content satisfies its own hypothesis. It is a warning for everything downstream: the natural next step in this program — putting a realistic quark-hosting representation into the AHMN machinery — silently loses a symmetry that the machinery's conventions assume. We paid this ourselves. An earlier version of our own naturalness ranking was computed on the half-domain, and its minimizer pinned to the $\alpha_2 = \frac{1}{2}$ boundary for five of nine representations; we diagnosed it as an implementation defect and fixed it numerically. It was not a defect. It was an inherited symmetry that our representations do not have.

Established here: the half-domain reduction inverts on the physical class, verified on the potential itself. *Left open:* what the admissible class gets *instead* — the subject of the next section, and not what one would guess.

6 What the notch buys: an exact decoupling

The odd-m pairing buys periodicity. The even-m notch buys something quite different, and the difference is the reason a naturalness ranking in this class has been so hard to stabilize.

The shape of what follows is familiar from a different context. Appelquist and Carazzone proved that heavy states drop out of low-energy observables up to suppressed corrections [31], and the lesson absorbed since is double-edged: whatever is protected from an observable is also *invisible* to it. A quantity that a symmetry forbids from contributing cannot be measured by what it does not contribute to. What we find below is a discrete, exact version of that trade. The notch does not suppress the fermion representation's contribution to a particular observable; it annihilates it. And the price is exactly the one Appelquist and Carazzone's readers learned to expect: that observable can no longer tell one representation from another.

The leading Kaluza–Klein order of the α_2-curvature at the half point is built *entirely* out of the even-m imbalance $\delta(m) := \mathrm{mult}(m,0) - \mathrm{mult}(m,1)$:

$$\mathcal{D} \;=\; \pi^2\Big(\sum_{m\equiv 2\,(4)} m^2\delta(m) \;-\; \sum_{m\equiv 0\,(4)} m^2\delta(m)\Big). \tag{9}$$

Nothing about (9) is new, and it is worth being precise about that before drawing the consequence. It is the $(k_1,k_2) = (0,\pm 1)$ term of the standard Coleman–Weinberg lattice sum differentiated twice at $\alpha_2 = \frac{1}{2}$, where the phase $\cos\big(\pi k_2(\frac{m}{2}+q)\big)$ kills odd m and leaves $(-1)^{q+m/2}$ on the rest; the manoeuvre of substituting the special Wilson-line value, watching a residue class cancel and reading the sign is written out in Hosotani, Noda and Takenaga [39, Eqs. (6.5)–(6.7)] for this very orbifold and in Haba, Hosotani, Kawamura and Yamashita [40, Eq. (46)]. Nor is the modulus-four grading as structural as it looks: substituting $m = 2m'$ collapses it to a plain alternating weight,

$$\mathcal{D} = -4\pi^2 \sum_{m'} (-1)^{m'} m'^2\, \delta(2m'),$$

so (9) is the standard leading-winding curvature written in residue classes and should be read as such. What does not follow from any of that is the corollary, because the notch is precisely the statement $\delta(m) = 0$ for all even m. Therefore:

Corollary 3 (Decoupling)**.** *For every Standard-Model-admissible representation, the fermion sector's contribution to $\mathcal{D}$ vanishes identically. $\mathcal{D}$ reduces to the gauge sector alone — the same number for every admissible representation, independent of (a,b,c).*

Numerically, $\mathcal{D} = -27.6$ for all eight admissible representations, identical to the digit, while the non-admissible **35**, which breaks the notch, sits apart at -185.5. This is the familiar shape of a conservation law: what makes the notch a clean *detector* of admissibility makes it blind as a *discriminator* among admissible representations.

A refuted inference, recorded deliberately. It is tempting to conclude from $\mathcal{D} < 0$ uniformly that $\alpha_2 = \frac{1}{2}$ is a curvature ridge for every admissible representation. That inference is false, and we record it because the failure is instructive. Against the true curvature (converged, $K_{\max} = 16$):

rep	**35**	**60**	**84**	**140***a*	**140***b*	**224**	**280**	**360**	**756**
$\mathcal{D}$ (leading order)	-185.5	-27.6	-27.6	-27.6	-27.6	-27.6	-27.6	-27.6	-27.6
true $\partial^2_{\alpha_2} V$ at $\frac{1}{2}$	-200.6	$+21.2$	-5.8	-85.1	-188.6	$+16.8$	$+255.2$	-160.4	$+245.5$
sign agrees	yes	**no**	yes	yes	yes	**no**	**no**	yes	**no**

The leading-order proxy mispredicts the sign in four of eight. It is right on the **35** — the one representation where the fermion sector dominates it, and the case against which such a proxy is implicitly calibrated. It was validated in the only regime where it works and then applied in the regime where the notch has removed its content.

The consequence that matters. Putting the two together:

> The notch does not fix the sign of the α_2-curvature; it *annihilates the leading term.* For every admissible representation the leading Kaluza–Klein order is degenerate — identical, gauge-only. Everything that distinguishes one admissible representation from another lives in the subleading terms.

This explains from theory, rather than from post-mortem, why a naturalness ordering over this class is delicate: the candidates are indistinguishable at leading order *by a theorem*, and the quantity that separates them is exactly the one a coarse minimization resolves worst. Together with Corollary 2, the two known failure modes of such a computation are consequences of the same parity flip.

Established here: an exact decoupling, and the structural origin of the fragility. *Left open:* the subleading structure itself — which now has a target, since the leading order is known in closed form to be a constant.

7 Measuring the space instead of the point

Every fragility above comes from evaluating at a point: a minimizer pinned to a boundary, a curvature read at that boundary, a Hessian taken by finite differences, a near-zero eigenvalue in a flat direction. The one exact statement of this paper — the notch — is a statement about the full Fourier content of V over the torus. That asymmetry suggests changing instrument.

For a quadratic well the sublevel set $S(\varepsilon) = \{V < V_{\min} + \varepsilon\Delta V\}$ is an ellipse with semi-axes $\sqrt{2\varepsilon\Delta V/\lambda_i}$, and a uniform ellipse has second-moment tensor $M_{ii} = a_i^2/4$. Both curvature eigenvalues therefore follow from an *integral* over the basin,

$$\lambda_i = \frac{\varepsilon\,\Delta V}{2\,M_{ii}}, \tag{10}$$

with no derivatives and no step size. Two corrections are required and both were found by testing rather than by reasoning: ε must lie in the quadratic regime (at $\varepsilon = 0.15$ the sublevel set covers 40–76% of the torus and M saturates against the torus itself, reading a genuine 413:1 anisotropy as 1.17:1), and $S(\varepsilon)$ must be restricted to the connected component of the global minimum, since the torus carries symmetry-equivalent minima and an unrestricted set measures the separation *between* minima — the signature being $\lambda \propto \varepsilon$ instead of a plateau.

Corrected, the measure has a plateau at $\varepsilon \in [3 \cdot 10^{-4}, 10^{-3}]$ reproducing the finite-difference Hessian for eight of nine representations. On the ninth — $(0, 3, 1)$, the one where finite differences are untrustworthy — the drift under $K_{\max}\colon 4 \to 14$ falls from 68.9% to 0.6%. Two facts established along the way are worth recording: the vacuum *location* is already converged at $K_{\max} = 4$ (identical to four decimals up to $K_{\max} = 14$; the **35** lands at $(0.4360, 0.2986)$ against AHMN's quoted $(0.438, 0.299)$, confirming the calibration), so under-convergence never affected where the vacuum is, only the curvature there; and the low mode of $(0, 3, 1)$ is genuinely *anharmonic*, with no ε-plateau at all, so a quadratic mass is simply the wrong descriptor for that direction rather than a hard-to-measure one.

Established here: an integral replacement for the Hessian, validated where the Hessian is trustworthy and stable where it is not. *Left open:* the subleading ranking itself.

8 A preflight, and who should run it

Part I ended by packaging its Existence theorem as a reusable oracle rather than leaving it as a claim about one model. The same is worth doing here, because Corollary 2 is not a fact about $(3, \mathbf{60})$ — it is a check that anybody building on this class of construction should run *before* spending computer time.

What it is for. The reduction (8) is the kind of step that travels as a convention. It appears once, is justified once for one field content, and is thereafter inherited — reasonably, since re-deriving every fundamental domain would make the literature unreadable. The trouble is that its hypothesis is representation-dependent and inverts on exactly the representations one would want to use next. A minimization run on the halved region for such a representation does not fail loudly; it returns a boundary point that looks like a vacuum. That failure mode is silent, which is why it deserves a tool rather than a remark.

What it does. `notch_preflight.py` takes either Dynkin labels or an explicit (m, q) multiplicity table and reports which pairing the tower satisfies, which fundamental region is therefore legal, and whether the representation can host a quark cell at all. On labels it applies Corollary 1 and Theorem 1 directly; on a table it *measures* both pairings rather than assuming either, so it remains valid for spectra outside the $SU(4)$ setting where our theorem applies. Typical output:

`notch_preflight.py 0 2 1`	`notch_preflight.py 4 0 0`
centre charge: odd	centre charge: even
tower pairs at: even m	tower pairs at: odd m
odd-m hypothesis: **FAILS**	odd-m hypothesis: **HOLDS**
legal region: $\alpha_2 \in [0, 1]$ **(full torus)**	legal region: $\alpha_2 \in [0, \frac{1}{2}]$
hosts a quark cell: **yes**	hosts a quark cell: no

Who should run it. Anyone extending a gauge–Higgs orbifold model to a new matter representation, before minimizing; anyone reusing a published fundamental domain with different field content; and anyone who finds a minimizer settling repeatedly on a domain boundary — which, as §5 records, is what a missing symmetry looks like from the inside. The check costs nothing: it is one parity bit.

Established here: the paper's operative consequence as a runnable check. *Left open:* the same question for other orbifold groups, where the alphabet in (5) changes and the dichotomy must be re-derived rather than reused — exactly the mistake this tool exists to prevent.

9 Ledger

As in Part II we set out ownership item by item rather than in prose; Table 4 is that accounting.

Where we looked. For a general criterion deciding when a $GL(n)$ character vanishes identically on $O(n)^-$ — of which Theorem 1 would be a corollary — we read Koike–Terada [18], King [15], Littlewood [17], Fauser–Jarvis–King–Wybourne [19], Ayyer–Kumari [20], Kumari [21], Ayyer–Behrend [22] and Albion [24], and found none. Koike–Terada cannot contain one: they define $R(O(n))$ as a subring of $R(SO(n))$, and since V_μ and $V_\mu \otimes \det$ restrict to the same $SO(n)$-irreducible the associate involution is quotiented out before the question can be posed. One near neighbour deserves explicit separation: Ayyer–Kumari's Example 2.16 states an $O(4)$ parity criterion at $(x, -x)$, which looks identical to ours and is not — their torus element has determinant $+1$, the *proper* component.

Updated in v4. Littlewood's book, Ch. XI, has since been obtained and read: it contains the self-associate mechanism (§11.6, pp. 226–227; §11.8, p. 235), so we cite it for that, but it does not contain the criterion on the $GL(n)$ weight λ. King's 1971 paper in its own exposition remains unreachable and is still flagged as unread; its abstract and a faithful reproduction elsewhere indicate it supplies the tool, not the theorem. And the search recorded above was itself deficient: it was run with a literature tool whose query builder silently discarded the discriminating terms of every query, so a null result from it certified nothing. The re-run, with that fixed and with INSPIRE-HEP and zbMATH Open added, is what produced the citations to [6, 7, 23, 5] that this version adds.

Item	Inherited (cite)	Ours
centre-charge congruence	$\mathbb{Z}_4$ n-ality; fractional Y locked to the centre [9, 10, 11, 12, 8]; congruency class as the invariant controlling Wilson-line breaking, in the Hosotani model itself [6, 7]	nothing — as in Part II, we claim none of it
orbifold, twists (m, q), one-loop potential, Eq. (8)	AHMN [4, 5]; the $\{A+B(-1)^{k_2}\}$ coefficients for the **35** are [5, Eq. (3.25)]	the derivation of (m, q) from the boundary conditions, and the multiplet reading of AHMN's Table 1
parity lemma 1	elementary (symmetry + homogeneity of s_λ)	presented as a remark, not a result
$D \equiv 0$ mechanism	$O(2n)$ associate/det-twist; Littlewood restriction; King modification rules [14, 16, 13, 22]	the explicit arithmetic classification (Thm. 1) and its Lean certificate
the dichotomy, Cor. 1	—	ours
inversion of (8), Cor. 2	—	ours, and verified on the potential directly
decoupling, Cor. 3	—	ours

Table 4: What each result owes. The novelty of §4 is a statement, not a phenomenon and not a technique; the novelty of §§5–6 is the physics.

10 Reflections

A null from a search tool is not a gap (v4). The previous version of this section recorded, correctly, that we had looked for a general non-identity-component criterion and not found one. What we did not record — because we did not know it — is that the literature tool used for part of that search was defective in a way that could only ever produce silence: its query builder kept the first few tokens of a query by position rather than by specificity, so the discriminating terms were discarded before the search ran, and several sources were not installed at all, leaving everything published before 1991 invisible. A tool in that state returns few or no references, and few or no references reads exactly like novelty. Two of the citations added in this version, [6, 7], are from 1988 and 1989, carry over one hundred citations each, and appear as the first two results of the first query once the tool works. We record this not as an apology but as a methodological point with teeth: *a negative search result is evidence only in proportion to the demonstrated sensitivity of the search*, and the sensitivity must be established on a query whose answer is already known before any null is believed. The same discipline this paper applies to physical inferences — see the next paragraph, and §6 — applies to bibliographic ones, and we had not been applying it.

An inherited convention carries an unstated hypothesis. Equation (8) travels as a convention — a fundamental domain, the kind of thing one copies without re-deriving. Its justifying hypothesis does not travel with it, and in this case the hypothesis inverts exactly on the class one would want to use next. We suspect this is not an isolated case in orbifold model building, where fundamental domains, Z_2 assignments and "without loss of generality" reductions are routinely inherited across papers with different field content. The cheap protection is to carry the hypothesis as an executable check rather than as a remark.

The centre as a constraint on the potential, not on the matter. The shift in viewpoint is easy to lose among the corollaries, so we state it bluntly.

For fifty years the centre of the gauge group has been used in one direction. Dirac's quantization, 't Hooft's phases, Hucks and Saller's reading of the Standard Model's own $\mathbb{Z}_6$: in each the centre is a condition on *which representations may exist*, or carry which charges. Part II used it that way too, and said so. The centre restricts matter.

What we have found is that in a gauge–Higgs construction the same invariant also restricts the *potential*. Not by making some contribution small: by removing an entire Fourier sector of the Hosotani potential identically, for every representation of the wrong parity. The Higgs in these models is a Wilson line; a period of Wilson line is a central gauge transformation; so the centre acts on the Higgs sector as directly as it acts on the matter sector, and its charge decides what can appear there. Figure 5 is that statement drawn: for eight of nine representations an arithmetically specified sublattice of harmonics is not small, it is empty.

We would put it this way. The centre charge has been read as a gate on matter. It is also a gate on the potential, and the two gates open in opposite directions — which is why a condition that lets a representation host the Standard Model is the same condition that denies its Wilson-line potential a symmetry the literature has been assuming.

A conjecture, stated so it can be killed. The mechanism we can prove is $SU(4)$-specific in its arithmetic but not in its logic, and it is more useful to say what we expect than to stop at what we computed.

> *Conjecture (centre-parity selection, general form).* Let an orbifold gauge–Higgs model have a Kaluza–Klein spectrum whose twist labels are captured by a character χ_λ evaluated on an alphabet A, in the manner of (5). Suppose a unit shift of a Wilson-line modulus carries A to $\zeta \cdot A$ up to a Weyl element, for some ζ in the centre of the gauge group, of order d. Then the harmonics of the one-loop potential split into d sectors labelled by the centre charge modulo d, and the sectors on which $\zeta^{\,c(R)} \neq 1$ are identically absent from the potential of a representation R.

We prove the case $d = 2$, in the model where every ingredient can be computed exactly, and we verify the structural precondition across $SU(N)$ below. What we have not done is exhibit a second model, of different d, in which the sectors appear as predicted; that is the obvious test and the obvious way to kill this. A $\mathbb{Z}_N$ orbifold with $N > 2$, or a Wilson line whose unit shift realizes a central element of order three, would settle it in one calculation.

Is this an accident of $SU(4)$? It is the first question to ask of (7), and it has a sharp answer rather than a hopeful one.

The principle needs the shift to equal a central element *up to Weyl*, which is the multiset condition $-A = A|_{t\to -t}$ on the alphabet $A = \{\varepsilon_i\, t^{c_i}\}$, with $\varepsilon_i = \pm 1$ from the $\mathbb{Z}_2$ twist and c_i the Wilson-line charge. Split the slots by c_i. Those with $c_i = \pm 1$ satisfy the condition automatically, since $-\varepsilon t^{\pm 1}$ and $\varepsilon(-1)^{\pm 1} t^{\pm 1}$ are the same monomial. Those with $c_i = 0$ do not: negation flips their sign and nothing in t can undo it, so they must pair off, a $+1$ against a -1. *The Wilson-line-neutral slots must carry balanced twist signs*, and in particular there must be an even number of them.

For a single Wilson-line direction that means $N - 2$ even, hence N even. A brute-force sweep over every alphabet of this shape confirms it:

$SU(N)$	N=3	N=4	N=5	N=6	N=7	N=8
alphabets admitting the rule	0	192	0	4880	0	134400
neutral slots when it holds	—	2	—	4	—	6
verdict	none	**exists**	none	**exists**	none	**exists**

So the mechanism is neither an accident nor universal: it lives on the even unitary groups, and $SU(4)$ **is the smallest group that can carry it**. The odd groups are excluded for a reason

one can state in a sentence — an odd number of spectator slots cannot pair its twist signs — and that reason is structural, not numerical.

Two honest limits on the sweep. It assumes the AHMN shape: one $\mathbb{Z}_2$ twist contributing signs, and one Wilson-line direction with charges in $\{0, \pm 1\}$. A $\mathbb{Z}_N$ orbifold, or a Wilson line with higher charges, changes the condition — the relevant central element then need not have order two, and the congruence that appears would be modulo its order rather than modulo 2. We have not swept those, and we do not claim them.

The method is a shadow. It is worth naming what was actually done, because the same move is available elsewhere. Nothing in this paper studies an $SU(4)$ representation directly. What it studies is a *shadow*: the specialization $D(t) = s_\lambda(1, -1, t, t^{-1})$ collapses a four-dimensional weight system onto a single Laurent polynomial, and the orbifold is what chooses the wall to project onto — the two $\mathbb{Z}_2$ parities and the Wilson line fix the alphabet, and we do not get to pick it. Read that way the results are statements about shadows. The dichotomy (Corollary 1) says the shadow is always symmetric, with the centre charge deciding about which axis. The notch says the shadow has holes, at arithmetically exact positions — and Figure 5 is those holes. The degenerate class $\mathcal{Z}$ of Theorem 1 is the strangest of the three: those representations cast *no shadow at all* on this wall. They are not small — $\mathcal{Z}$ contains representations of every size — and from here they are simply invisible, which is why the physics never sees them and why they had to be classified by algebra rather than found by inspection.

The practical moral is that the wall is chosen for you. A different orbifold projection is a different alphabet in (5), hence a different shadow and a different dichotomy; that is precisely why §8 warns against reusing the result rather than re-deriving it. What generalizes is not the answer but the question: given a compactification, ask what specialization it imposes on the character, and read the physics off the shadow.

A gate can be a filter. The picture that emerges is that the representation acts as a *filter* on the vacuum wave: $V(\alpha)$ is a Fourier series in the Wilson line, the Kaluza–Klein sum is a Fourier transform on the hidden torus, and the multiplicity spectrum of the representation decides which components pass. Corollary 3 says that admissible representations are precisely those whose filter is blind at leading order — selecting a representation that can host the Standard Model automatically installs a notch in its own Higgs potential. Whether that blindness is a nuisance or a design principle (choose the representation to shape the filter) is the question we would most like to answer next.

The leading order is a constant, so aim below it. Because Corollary 3 fixes the leading term in closed form, the naturalness problem is now well posed rather than merely hard: compute the subleading structure with an analytic gradient and explicit convergence control, against a leading order known to be degenerate. That is a sharper task than making a minimizer more careful, and it is the first time there has been a reason to expect it to be necessary.

11 Conclusion

Part II identified its first gate as classical and gave it away. This paper shows the gate was doing more work than the sector it was found in, and the reason is one sentence: advancing the Wilson line by one period is, up to a Weyl reflection, multiplication by the central element $-\mathbf{1}$, so the centre acts on the Higgs sector as directly as it acts on the matter sector. A representation answers with a single bit, and that bit decides which harmonics of the Hosotani potential are permitted to exist at all. What has been read for fifty years as a restriction on *which representations may carry which charges* is also a restriction on *what the potential may contain*.

Concretely: the parity of the $\mathbb{Z}_4$ centre charge — equivalently $a + c$ — decides which of two Kaluza–Klein twist pairings a representation satisfies, and the two are mutually exclusive outside an explicitly classified degenerate class. One of those pairings is the hypothesis under which the two-Higgs-doublet gauge–Higgs literature halves its Wilson-line fundamental domain; the other is what every Standard-Model-admissible representation has instead. The consequence is concrete and, we believe, unavoidable for the program: the half-domain reduction does not transfer, the full torus must be scanned, and the leading Kaluza–Klein order of the curvature at the half point is representation-independent, so any ordering of candidates by naturalness is a statement about subleading terms whether or not it is presented as one.

We claim none of the centre-charge congruence, exactly as before. What we claim is that it is not only a gate on matter: it is a gate on the potential, and the two gates open in opposite directions. From one witness, to the law of eligibility, to the discovery that the law's most classical clause governs the sector it was never asked about — that is the arc of the three papers.

Data availability and reproducibility. All scripts, weight data, Lean 4 sources and figure-generation code are provided as ancillary files; no data beyond what they generate underlies any claim here. Every number regenerates from the ancillary scripts: the bialternant sweep for Theorem 1, the Lean 4 source of Table 1, the (m, q) derivation and multiplet-spreading check of §2, the direct periodicity test of §5, the decoupling and curvature comparison of §6, and the basin measure of §7 with its ε and $K_{\max}$ sweeps.

References

[1] C. Marín, *Anomaly- and Tadpole-Compatible Fermion Completion of Six-Dimensional $SU(4)$ Gauge–Higgs Unification* (Part I), Zenodo (2026), DOI: 10.5281/zenodo.21429144.

[2] C. Marín, *Three Gates to a Quark Generation* (Part II), Zenodo (2026), DOI: 10.5281/zenodo.21432628.

[3] C.A. Scrucca, M. Serone, L. Silvestrini, A. Wulzer, *Gauge-Higgs Unification in Orbifold Models*, JHEP **02** (2004) 049, hep-th/0312267.

[4] K. Akamatsu, T. Hirose, N. Maru, A. Nago, *Two Higgs Doublet Model from Gauge–Higgs Unification in Six Dimensions*, arXiv:2603.05857.

[5] K. Akamatsu, T. Hirose, N. Maru, A. Nago, *Electroweak Symmetry Breaking in Two Higgs Doublet Model from 6D Gauge–Higgs Unification on $T^2/\mathbb{Z}_2$*, arXiv:2312.08608.

[6] A.T. Davies, A. McLachlan, *Gauge group breaking by Wilson loops*, Phys. Lett. B **200** (1988) 305.

[7] A.T. Davies, A. McLachlan, *Congruency class effects in the Hosotani model*, Nucl. Phys. B **317** (1989) 237.

[8] R. Slansky, *Group theory for unified model building*, Phys. Rept. **79** (1981) 1.

[9] J. Hucks, *Global structure of the standard model, anomalies, and charge quantization*, Phys. Rev. D **43** (1991) 2709.

[10] M. Saller, *The Central Correlations of Hypercharge, Isospin, Colour and Chirality in the Standard Model*, Nuovo Cim. A **111** (1998) 1375, arXiv:hep-th/9802112.

[11] D. Tong, *Line Operators in the Standard Model*, JHEP **07** (2017) 104, arXiv:1705.01853.

[12] R. Alonso, D. Dimakou, J. West, *Fractional-charge hadrons and leptons to tell the Standard Model group apart*, arXiv:2404.03438.

[13] A. Solymos, D. Jakab, Z. Zimborás, *Extendibility of Brauer states*, arXiv:2411.04597. (Prop. 1 and Thms. 5, 9 of App. A give, in modern notation, the Littlewood restriction rule, the $O(2k)$ self-associate fact, and the Newell–King modification rules.)

[14] R.C. King, *Modification Rules and Products of Irreducible Representations of the Unitary, Orthogonal, and Symplectic Groups*, J. Math. Phys. **12** (1971) 1588.

[15] R.C. King, *Branching rules for classical Lie groups using tensor and spinor methods*, Can. J. Math. **23** (1971) 176.

[16] M.J. Newell, *Modification rules for the orthogonal and symplectic groups*, Proc. Roy. Irish Acad. **54** (1951) 153.

[17] D.E. Littlewood, *Products and plethysms of characters with orthogonal, symplectic and symmetric groups*, Can. J. Math. **10** (1958) 17.

[18] K. Koike, I. Terada, *Young-diagrammatic methods for the representation theory of the classical groups of type B_n, C_n, D_n*, J. Algebra **107** (1987) 466.

[19] B. Fauser, P.D. Jarvis, R.C. King, B.G. Wybourne, *New branching rules induced by plethysm*, J. Phys. A **39** (2006) 2611, arXiv:math-ph/0505037.

[20] A. Ayyer, N. Kumari, *Factorization of classical characters twisted by roots of unity*, J. Algebra **609** (2022) 437, arXiv:2109.11310.

[21] N. Kumari, *Factorization of classical characters twisted by roots of unity: II*, arXiv:2212.12477.

[22] A. Ayyer, R.E. Behrend, *Factorization theorems for classical group characters, with applications to alternating sign matrices and plane partitions*, J. Combin. Theory Ser. A **165** (2019) 78, arXiv:1804.04514.

[23] S. Karmakar, *Character values at elements of order 2*, arXiv:2412.17324.

[24] S.P. Albion, *Universal characters twisted by roots of unity*, arXiv:2212.07343.

[25] G. 't Hooft, *On the phase transition towards permanent quark confinement*, Nucl. Phys. B **138** (1978) 1.

[26] J. Scherk, J.H. Schwarz, *Spontaneous breaking of supersymmetry through dimensional reduction*, Phys. Lett. B **82** (1979) 60.

[27] N.S. Manton, *A new six-dimensional approach to the Weinberg–Salam model*, Nucl. Phys. B **158** (1979) 141.

[28] D.B. Fairlie, *Higgs' fields and the determination of the Weinberg angle*, Phys. Lett. B **82** (1979) 97.

[29] Y. Hosotani, *Dynamical mass generation by compact extra dimensions*, Phys. Lett. B **126** (1983) 309.

[30] Y. Kawamura, *Triplet-doublet splitting, proton stability and extra dimension*, Prog. Theor. Phys. **105** (2001) 999, arXiv:hep-ph/0012125.

[31] T. Appelquist, J. Carazzone, *Infrared singularities and massive fields*, Phys. Rev. D **11** (1975) 2856.

[32] H. Weyl, *The Classical Groups: Their Invariants and Representations*, Princeton University Press (1939).

[33] D.E. Littlewood, *The Theory of Group Characters and Matrix Representations of Groups*, 2nd ed., Oxford University Press (1950).

[34] F.D. Murnaghan, *The Theory of Group Representations*, Johns Hopkins Press (1938).

[35] R. Howe, E.-C. Tan, J.F. Willenbring, *Stable branching rules for classical symmetric pairs*, Trans. Amer. Math. Soc. **357** (2005) 1601, arXiv:math/0311159.

[36] M. Ciucu, C. Krattenthaler, *A factorization theorem for classical group characters, with applications to plane partitions and rhombus tilings*, arXiv:0812.1251.

[37] N. Haba, Y. Hosotani, Y. Kawamura, *Classification and dynamics of equivalence classes in $SU(N)$ gauge theory on the orbifold S^1/Z_2*, Prog. Theor. Phys. **111** (2004) 265, arXiv:hep-ph/0309088.

[38] N. Haba, K. Takenaga, T. Yamashita, *A general formula of the effective potential in 5D $SU(N)$ gauge theory on orbifold*, arXiv:hep-ph/0401185.

[39] Y. Hosotani, S. Noda, K. Takenaga, *Dynamical gauge symmetry breaking and mass generation on the orbifold T^2/Z_2*, arXiv:hep-ph/0403106.

[40] N. Haba, Y. Hosotani, Y. Kawamura, T. Yamashita, *Dynamical symmetry breaking in gauge-Higgs unification on orbifold*, arXiv:hep-ph/0401183.

[41] N. Weiss, *Effective potential for the order parameter of gauge theories at finite temperature*, Phys. Rev. D **24** (1981) 475.

[42] D. J. Gross, R. D. Pisarski, L. G. Yaffe, *QCD and instantons at finite temperature*, Rev. Mod. Phys. **53** (1981) 43.

[43] A. Roberge, N. Weiss, *Gauge theories with imaginary chemical potential and the phases of QCD*, Nucl. Phys. B **275** (1986) 734.

[44] M. Ünsal, L. G. Yaffe, *Center-stabilized Yang–Mills theory: confinement and large N volume independence*, Phys. Rev. D **78** (2008) 065035, arXiv:0803.0344.

[45] Y. Kawamura, T. Miura, *Equivalence Classes of Boundary Conditions in $SU(N)$ Gauge Theory on 2-dimensional Orbifolds*, arXiv:0905.4123.

[46] K. Takeuchi, T. Inagaki, *A new classification method for equivalence classes on S^1/Z_2 and T^2/Z_3*, PTEP **2024**, arXiv:2401.09809.

[47] K. Takeuchi, T. Inagaki, *Trace conservation laws in T^2/Z_m orbifold gauge theories*, arXiv:2404.19411.

[48] F. Feruglio, *Are neutrino masses modular forms?*, in *From My Vast Repertoire: Guido Altarelli's Legacy*, World Scientific (2019) 227, arXiv:1706.08749.

[49] H.P. Nilles, S. Ramos-Sánchez, P.K.S. Vaudrevange, *Eclectic flavor groups*, JHEP **02** (2020) 045, arXiv:2001.01736; see also M. Ratz et al. for the orbifold-based development of modular flavour symmetry.